Report on the Proposed National Do-Not-Email Registry

Prepared for the Federal Trade Commission

Edward W. Felten
Professor of Computer Science
Princeton University[1]
May 2, 2004

1. Introduction

This report presents my analysis of the feasibility and efficacy of establishing a national Do-Not-Email Registry (DNER). It was prepared at the request of the Federal Trade Commission (FTC), and is based on a Request for Information issued by the FTC, and on the replies to that Request.

The rationale for a DNER is similar to that for the national Do Not Call list: to protect consumers from a rising tide of unwanted, and sometimes offensive, unsolicited commercial email. While the Do Not Call list is widely seen as a success, important differences between email and telephony, including the greater anonymity of email addresses compared to telephone numbers, and the lower cost of sending email messages, make the success of a DNER far from assured.

In this report, I will use the term "spam" to refer to commercial bulk email messages, other than legitimate messages such as those associated with opt-in mailing lists, transactional messages, and so on.

This report is structured as follows. First, I review a few relevant aspects of spammers' behavior. Next, I review the DNER technologies suggested in the responses to the FTC's Request, and discuss the security of the two types of technologies that were suggested. I go on to discuss some important issues related to the enforceability of the DNER's rules, and I conclude with some specific recommendations.

2. How Spammers Operate

The starting point for analyzing a DNER is to understand the two main technical challenges faced by a would-be spammer: how to gather a list of email addresses to which spam messages will be sent, and how to send messages that will get through to the people at those addresses. A successful DNER would operate by making one or both of

[1] Affiliation is given for identification purposes only. This report is not sponsored or endorsed by Princeton University.

these tasks more difficult. A poorly designed DNER might actually make the spam problem worse by helping spammers accomplish these tasks.

Gathering a mailing list of addresses to spam can be challenging because there is no central list of valid email addresses. Addresses can be "harvested" from web pages or the archives of mailing lists, they can be purchased from direct marketers, or they can simply be guessed. In practice, a spammer's mailing list may contain many obsolete or invalid addresses, and it can be difficult for the spammer to "purify" his list by weeding out the bad addresses.

A carelessly designed DNER might actually make the spam problem worse by helping spammers construct larger and more accurate mailing lists. To prevent this, we must avoid any design that gives bulk emailers unprotected access to the DNE list, because an unscrupulous spammer who received the list would simply send spam email to the people on the list. We must also take care to avoid any DNER design that helps spammers purify their mailing lists in other ways.

The second problem faced by a spammer is how to get messages through to the people on his mailing list. The obvious way to do this is for the spammer to buy a high-bandwidth connection to the Internet over which he can stream out messages rapidly from his computers. However, this approach tends to draw attention to the spammer, which often leads to the spammer's Internet Service Provider (ISP) cutting off service, or to the spammer's computers ending up on various anti-spam "blacklists."

Because of these hurdles, many spammers employ alternative approaches based on illegal practices such as fraud and electronic break-ins. For example, the spammer may break in to the computers of ordinary, unsuspecting Internet users and install stealthy programs, known as "zombies" or "bots," that cause the victims' machines to send spam. The spammer may also insert false and misleading information into the headers of the spam messages, making it very difficult to distinguish the spam messages from legitimate email or to find the true source of the messages. I will use the term "outlaw spammers" to refer to spammers who are willing to resort to such tactics.

To be effective, a DNER approach must make at least one of the spammers' two problems more difficult. At the very least, it must avoid making either of them easier. Below, I will evaluate the proposed DNER schemes according to these criteria.

3. Proposed Do-Not-Email Technologies

The replies to the FTC's Request differed in some important respects, but there were also broad similarities between them. I will recount the similarities before turning to the areas of difference.

The proposals would all rely on generally similar methods by which people could add their email addresses to the DNER. These are roughly similar to the registration approach used for the Do Not Call list, with some obvious extensions to account for the

fact that all requesters can be assumed to have email access. Because of these similarities, and because the methods suggested seem to be well designed, not much needs to be said about the registration process.

Most of the proposals would allow the owner of an Internet domain to register the entire domain for the DNER. (The domain is the part of an email address after the '@' symbol, e.g. aol.com, or ftc.gov.) Per-domain registration is clearly a good idea, since it increases user convenience, and reduces the size of the DNER (by storing one registration for an entire domain, rather than individual registrations for each address in the domain). If a DNER is adopted, it should include both per-address and per-domain registrations.

The most challenging part of the DNER design is the interaction with bulk emailers, because some bulk emailers will be legitimate and some will be spammers.
In their strategies for dealing with bulk emailers, the proposals fall into two general categories, one based on list scrubbing and the other based on the use of trusted remailers. In the list scrubbing approach, commercial bulk emailers would be given a way to remove addresses on the DNER from their mailing lists. In the remailing approach, commercial bulk emailers would be required to direct their bulk email traffic through a trusted remailing service, which would discard any messages that were addressed to parties on the DNE list, but would forward all other messages to their destinations.

List scrubbing is the cheaper of the two approaches. However, in my opinion, the cost difference between the two approaches is relatively small, so the decision to favor one over the other should be made on the basis of efficacy rather than cost.

I will now address the security implications of each of these approaches.

4. Analysis of List Scrubbing

The list scrubbing approach operates by allowing a bulk emailer to eliminate addresses on the DNER from its mailing list, so that email messages sent to addresses on the scrubbed list could never reach an address on the DNER.

List scrubbing can be implemented in several ways. The emailer might upload its list to the DNER administrator, which would scrub the list and return the scrubbed list to the emailer. Alternatively, the DNER administrator might provide a service that emailers could query in real-time to determine whether a given address is on the DNER. Finally, the DNER administrator could utilize cryptography, by giving emailers a list containing the cryptographic hash of each address on the DNER, thereby allowing emailers to check whether any given address is on the list, without directly disclosing the list to them[2].

[2] Cryptographic hashing can be thought of as a method for "anonymizing" an address, so that the original address cannot be recovered from the anonymized version. Giving emailers only the anonymized version of the list, and not the original list itself, helps to protect the original list from becoming a source of new addresses for spammers. However, due to the mathematical properties of cryptographic hashes, it is still

List scrubbing would give law-abiding commercial emailers a reasonable way to comply with the DNER rules. But list scrubbing has a fatal flaw: spammers can use it to purify their mailing lists. A spammer can do this by submitting his mailing list to the scrubbing service, and then comparing the original, unscrubbed list to the scrubbed version, to determine which addresses were removed in the scrubbing process. The spammer will know that the removed addresses are on the DNER list, and from this he can infer that they are valid addresses that are being read by human beings. By this stratagem, a spammer can use a list scrubbing system to improve the quality of his mailing list, thereby increasing the effectiveness of his spamming.

In my view, this flaw is reason enough to reject the list scrubbing approach. If implemented, this approach is likely to make the spam problem worse.

5. Analysis of Trusted Remailers

In the trusted remailer approach, the DNER list would be distributed only to a small set of trusted remailing services. Bulk commercial emailers would be required to route their bulk email through one of the trusted remailing services. The remailing services would discard any messages directed to addresses on the DNER list, and would forward all remaining messages on to their destinations.

Presumably there would be several trusted remailing services, and each would deploy a number of computers to provide remailing services. Because of the plurality of services, and because it would be easy for each service to expand its capacity simply by adding more computers, the remailing infrastructure could scale as necessary to handle the required number of messages. The cost of such an infrastructure could be quite low, on a per-message basis, and the remailing services would presumably charge bulk emailers a (very small) per-message fee to recoup that cost.

Like the list scrubbing approach, the trusted remailer approach allows law-abiding commercial bulk emailers to comply easily with the DNER requirements. Unlike the list scrubbing method, however, trusted remailers do not give spammers any way to purify their mailing lists. Emailers are told nothing about which addresses are or are not on the DNE list, and they are not told which of the messages they send through the remailers are discarded because their addressees are on the DNE list. Accordingly, information about the DNE list does not leak.

Because of this difference, the trusted remailer approach is clearly superior to the list scrubbing approach. If a DNER is to be implemented, I recommend that it be based on trusted remailers.

---

possible for a person who knows an email address to tell whether that address is on the anonymized list. So a system based on cryptographic hashes is roughly equivalent, from a security standpoint, to one that allows emailers to query a centralized database to check whether particular addresses are on the list.

6. Measures to Facilitate Enforcement

Regardless of which DNER approach is implemented, there are three measures that would help to facilitate enforcement of the DNER rules.

The first measure would require labeling of every commercial bulk email message that is subject to the DNER rules. The label might take the form of a header added to a message. (Email messages already contain many types of headers, most of which are read only by email-handling software and are not ordinarily displayed to users.) The idea would be to give the sender of a lawful message a way to state why that message was permitted – the label might say, for example, that the message was sent by an opt-in mailing list or service, or that the message was transactional in nature. Deliberate mislabeling of messages would presumably be forbidden.

The second measure would require the senders of commercial bulk email to register with the DNER administrator. Registration would make it easier to find emailers if enforcement action became necessary, and it would facilitate message authentication, as described below.

The third measure would establish a standard mechanism by which commercial bulk emailers could cryptographically authenticate their messages (for example, by digitally signing messages), so that recipients would know the true source of each bulk email message. This goes hand in hand with the registration of commercial bulk emailers – when an emailer registered, it could provide a copy of the public cryptographic key that it planned to use for authenticating its outgoing messages. The registration authority would then publish a list of the public keys of all bulk emailers in good standing; or, equivalently, the authority could give each such emailer a digital certificate attesting to its registration and its subsequent compliance.

If all three of these measures were put in place, then every legitimate commercial bulk email message would be both labeled and signed, so that recipients (and their spam-blocking filters) would be able to verify that a message came from a sender in good standing with the DNER administrator, and that that sender had asserted a rationale for the message's legality. Spam-blocking filters presumably would be programmed to accept such messages; and the system would foster accountability, since commercial bulk emailers would have to commit themselves to assertions about why their messages were allowable, and would risk having their registration revoked (and possibly facing other enforcement actions) if they failed to comply.

If any kind of DNER is implemented, I recommend that it be accompanied by these three measures – labeling of messages, registration of commercial bulk emailers, and authentication of commercial bulk email – to foster enforcement of the DNER rules. Indeed, these measures are useful even if no DNER is deployed.

7. Can DNER Rules Be Enforced?

Any DNER scheme, or for that matter any anti-spam law or mechanism, will face very serious enforceability problems. The difficulty is that outlaw spammers are likely to ignore the rules and send noncompliant spam messages disguised as ordinary email. Outlaw spammers have already shown themselves willing to break the law, for example by breaking into innocent users' machines, and a DNER is unlikely to increase significantly the risk of legal enforcement that outlaw spammers face. In addition, many outlaw spammers operate from overseas in order to frustrate enforcement.

For this reason, spam will continue to be a serious problem whether or not a DNER is implemented. The relevant question is whether the DNER will reduce the spam problem enough to justify the cost of its implementation.

The argument against a DNER is that it will not deter the outlaw spammers, who will continue to behave as before, flooding our email inboxes with unwanted messages that are difficult to distinguish from legitimate email. If, as seems likely, the core of the spam problem is caused by outlaw spammers, then a DNER may do little if anything to reduce the level of spam.

The best argument in favor of a DNER is indirect. A DNER, coupled with labeling and authentication mechanisms, will make legitimate commercial email (such as transactional messages and opt-in mailings) easier to recognize, since legitimate messages will be labeled and will come from authenticated sources. If spam filters can recognize these legitimate commercial messages and give them special treatment, then the filters can apply harsher scrutiny to any remaining messages that appear to be commercial; and that harsher scrutiny may reduce the number of outlaw spam messages that get through.

It is worth noting, however, that these benefits can be achieved without a DNER, simply by implementing the labeling, registration and authentication mechanisms described above in Section 6. These mechanisms, by themselves, will tend to increase the efficacy of privately operated spam filters, and thereby reduce the amount of spam that ends up in users' inboxes.

The best policy, in my opinion, is <u>not</u> to implement a DNER, but instead to focus on giving legitimate commercial emailers a way to distinguish their messages from the flood of spam, in order to increase the effectiveness of private spam filters while decreasing the number of legitimate messages that are incorrectly filtered. This can be done by requiring labeling of commercial bulk email messages, registration of commercial bulk emailers, and authentication of the sources of such email messages.

While a properly designed DNER will do little if any harm, it will do little good in addressing the spam problem.

8. Recommendations

To summarize, in this report I have offered six specific recommendations.

1. If a DNER is deployed, the procedures for adding an address should be similar to those used for adding a number to the Do Not Call list. Alternative procedures should exist for adding entire domains to the DNER.
2. The costs of the various DNER proposals are roughly comparable, and are low in comparison to the potential benefits of reducing spam, so any choice between DNER alternatives should be made based on efficacy rather than on cost.
3. DNER designs based on the list scrubbing approach should be rejected, because they will help spammers improve the quality of their mailing lists. Such designs will probably make the spam problem worse rather than better.
4. If a DNER is to be deployed, it should be of the trusted remailer type. Trusted remailer designs can be implemented at reasonable cost, and they lack the drawbacks of list scrubbing.
5. Deployment of a properly designed DNER will do little if any harm; but it will also fail to reduce the spam problem, because outlaw spammers will simply ignore the DNER rules. Therefore a DNER should not be deployed.
6. Regardless of whether a DNER is deployed, steps should be taken to ease the enforcement of anti-spam rules and to help email recipients distinguish legitimate commercial emails from spam messages. Specifically, it would be useful to require labeling of all commercial bulk email messages, with labels stating into which legal category each message falls; to register commercial bulk emailers; and to require or at least encourage the authentication of legitimate commercial bulk email messages.